

Does Nash Envy Immunity?

Ching-Hua Yu *

March 10, 2017

Abstract

The most popular stability notion in games should be Nash equilibrium under the rationality of players who maximize their own payoff individually. In contrast, in many scenarios, players can be (partly) irrational with some unpredictable factors. Hence a strategy profile can be more robust if it is resilient against certain irrational behaviors. In this paper, we propose a stability notion that is resilient against envy. A strategy profile is said to be *envy-proof* if by deviation, each player cannot gain a competitive edge with respect to the change in utility over the other players. Together with Nash equilibrium and another stability notion called immunity, we show how these separate notions are related to each other, whether they exist in games, and whether and when a strategy profile satisfying these notions can be efficiently found. We answer these questions by starting with the general two player game and extend the discussion for the approximate stability and for the corresponding fault-tolerance notions in multi-player games.

*Computer Science Department, University of Illinois, Urbana-Champaign, cyu17@illinois.edu.

Contents

1	Introduction	3
2	Related Works and Discussion	4
2.1	Terminology Disambiguation for Envy-Related Notions	4
2.2	Irrational Players in Games	5
3	Definition	6
4	Warm-up for Immunity and Envy-proofness	6
5	On (Non)-Existence	7
6	On Efficient Solvability	9
7	Envy-proof Nash and Immune Nash	9
7.1	Proof of Theorem 3	10
8	On Approximation	12
8.1	Proof of Theorem 5	13
9	Multiplayer Games	14
9.1	Preliminaries for Multiplayer Games	15
9.2	Coalitional Envy-proofness	15

1 Introduction

We study an inherent yet unpredictable factor of players in games that is related to the notorious envy¹ of humanity. It is observed that in a company, the (un)happiness of employees can depend on, instead of their own salary, the comparison between their salary and their colleagues'. In some scenarios, the situation can be worse. For example, a (somehow) irrational bidder in an auction can set a bid higher than the true value of an item just because he is unhappy to see his rival bidder win the item.

In general, the outcome of games can deviate from the predicted equilibria, and the players' behavior can disobey their presumed utility function. Even though one remedy is to redefine the utility function and make it as accurate as possible, sometimes the utility function of a real player is hard to predict or to accurately model, and due to the complex humanity, sometimes faulty players can just become spontaneous or unwilling to follow a reasonable utility under a model. Hence it is attractive if an equilibrium has more robust properties against these inherent factors and can prevent potential irrational behavior of players.

To capture the inherent irrational factor due to envy, we define a property called *envy-proofness*. An envious player would not be satisfied with her current state if she discovers that she can gain more utility "than others", or she would lose less utility "than others" by changing her strategy. An envy-proof strategy profile is stable with respect to this agitation. That is, A player cannot gain a competitive edge with respect to the utility change over the other players. Note that an envy-proof state may not be *envy-free*: A player in an envy-proof state can envy the other's utility compared to her own, but she has no incentive to deviate. Envy-proofness is independent of Nash equilibria in general; a Nash equilibrium can be envy-proof or not. Let's look at its definition for two-player games first.

Definition 1.1 (Envy-proof Profile). *In a two-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$, we say a strategy profile $\mathbf{x} = (\mathbf{x}_0, \mathbf{x}_1)$, a random variable over $A_0 \times A_1$, is envy-proof if and only if the following holds:*

$$\begin{aligned} & \forall b, \forall x'_b \in A_b, u_b(\mathbf{x}'_b : \mathbf{x}_b) - u_b(\mathbf{x}) \leq u_{\bar{b}}(\mathbf{x}'_b : \mathbf{x}_b) - u_{\bar{b}}(\mathbf{x}), \\ & \text{(where the notation } u_b(\mathbf{x}) = E_{x \sim \mathbf{x}}[u_b(x)]\text{),} \\ & \text{abbreviated as} \\ & \forall b, \Delta_b u_b(\mathbf{x}) \leq \Delta_b u_{\bar{b}}(\mathbf{x}) \end{aligned} \tag{1}$$

	movie	shopping
movie	4,4	1,3
shopping	3,1	3,3

Figure 1: A coordination game. (movie, movie) is a Nash equilibrium but not envy-proof, while (shopping, shopping) is an envy-proof Nash equilibrium.

As an example, consider the coordination game described in Figure 1. The two players like shopping no matter whether they do it alone or together. On the other hand, they only like to go to the movies together. While (movie, movie) and (shopping, shopping) both sound like a good coordinated strategy profile, one day, the two players quarrel with each other, behave irrationally, and even try to make each other unhappy regardless self-happiness. Then (movie, movie) is no longer a good strategy since it is unstable. It is unstable because each player has an incentive to change the strategy, and assuming the other's strategy is fixed, this deviation can cause the other player to lose more happiness (losing 3 points vs losing 1 point). On the other

¹In English, malicious envy is more like jealousy, but here we adopt a general term *envy* from the philosophy.

hand, (shopping, shopping) remains a good strategy because any unilateral deviation of a player can only cause his own loss.

In addition to envy-proofness, we study another stability notion called immunity. In [11, 13], Gradwohl and Reingold proposed and studied (ϵ, t) -immunity for special games in a fault-tolerance genre. Essentially a strategy profile is (ϵ, t) -immune if any t players cannot deviate and make some other player's expected utility decrease more than ϵ (in absolute value). They showed that in some special games (which we will review briefly later), every Nash equilibrium is (ϵ, t) -immune. While we will adapt the envy-proof notion to fault-tolerance issues in multi-layer games later in Section 9, we focus more on the fundamental properties such as the existence and the computational complexity of exact as well as approximate solutions in general games. Adapting the notion of immunity to two player games, we say a strategy profile is immune if each player's unilateral deviation cannot decrease the other player's utility.

Definition 1.2 (Immunity Profile). *In a two-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$, we say a strategy profile \mathbf{x} is immune if and only if:*

$$\forall b, \forall x'_b \in A_b, u_b(x'_b : \mathbf{x}_{-b}) \geq u_b(\mathbf{x}),$$

This paper is then centered around the three stability notions: Nash equilibria, envy-proofness, and immunity and developed from multiple dimensions. First we start by inspecting general two-player games and illustrate a clean connection between the three stability notions in Section 4. And in Section 5, we study the (non)existence of these notions. We show that while every two-player game has at least an immune strategy profile and an envy-proof strategy profile, an immune Nash equilibrium and an envy-proof Nash equilibrium do not always exist. Then following these observations regarding existence, in Section 6 we are concerned about whether they can be efficiently found in a game. We show that while finding an immune profile is *PPAD-complete*, finding an envy-proof strategy profile is polynomial-time solvable.

Based on the above discoveries, we check advanced properties of envy-proof Nash equilibria and immune Nash equilibria. In Section 7, somehow surprisingly, we show while determining the existence of envy-proof Nash equilibria is *NP-complete*, determining the existence of immune Nash equilibria is polynomial-time solvable. Next, because of the PPAD-hardness of finding an immune strategy profile and the NP-hardness of determining the existence of an envy-proof Nash equilibrium, in Section 8, we relax the definition of immunity and envy-proofness to resort to approximate solutions. Finally in Section 9, we discuss the multi-player fault-tolerance version of these notions. Particularly we show positive results of (ϵ, t) -coalitional envy-proofness for γ -varied games, γ -sensitive games, and anonymous games.

2 Related Works and Discussion

2.1 Terminology Disambiguation for Envy-Related Notions

Envy-related notions appeared a few times in several areas. In 1958, Gamow and Stern [9] introduce an envy-free concept for fair cake-cutting and chore division. Essentially the envy-free concept means that each partner believes that their share is at least as great as any other share. Note that the envy-freeness in the problem is quite different from our envy-proofness for general games. Particularly instead of comparing the utilities of the players, we want to compare the change in utilities of the players. That is, the utility between players can be substantially different but a strategy profile would be envy-proof if by any unilateral deviation of a player, the increase of her utility is less or the decrease of her utility is more than the others'.

The term envy-freeness and local envy-freeness also appeared in the problem of generalized second-price (GSP) auction, given by Varian [26] and Edelman et al. [8] independently. The concept there is to

capture the utility optimization of rational players (while excludes some plausible equilibria that cannot provide a good utility). In contrast, the envy-proofness in this paper is motivated by irrational players. Particularly an envy-free bid profile in GSP is necessarily an equilibrium, while an envy-proof profile in a game is not necessary a Nash equilibrium. For example, an irrational bidder in an auction can deviate to decrease his payoff if this can decrease his rival bidder's payoff more.

2.2 Irrational Players in Games

While the rationality of the players is a primal assumption in game theory, recently variant points of view are proposed to study or to explain irrational player behaviors in the computer science literature in recent years.

Fault Tolerance. One of the closest directions to our work is about fault tolerance of Nash equilibria against certain kinds of malicious attacks in large games. Following the line, Bernheim et al. [5], Kalai [16], Gradwohl and Reingold [12, 13] showed that Nash equilibria are robust against certain kind of attacks in general games satisfying specific properties. Particularly, Gradwohl and Reingold [13] showed that an honest player's payoff in a Nash equilibrium remains almost the same when the number of Byzantine faulty players is less than a threshold. The kind of tolerance is called immunity.

An instant question is that even though approximate immune Nash equilibria exist in games with a special form, it is unclear whether exact/approximate immune profiles or immune Nash equilibria exist in general games and whether it is computationally hard to find a solution. We answered these questions in this work, and so in some sense, we provide a fundamental inspections for the complement points following the line. Also, we study a new notion called envy-proofness that enriches the player's behaviors between fully rational and fully malicious. Interestingly, we will show Nash equilibria, immunity, and envy-proofness are different but relative notions in these questions.

Spitefulness. Another related direction is about the spiteful/altruistic setting. The setting says that a player's "true" utility is a combination of his own and the others. (See Chen[6] and Chen and Micali[16] for detailed reviews and references on this line.) One major difference between their direction and ours is that the spiteful/altruistic setting assumes certain knowledge of the external factors (called externalities) that decides the true utility function. This is a necessary step in their approaches since their goals are about how the profit of an auction problem or the social welfare of a game would compromise due to these externalities and how an (approximate) solution of an optimization problem according to the redefined utilities can be obtained.

In contrast, we consider scenarios whether the external factors are unpredictable and focus on issues about Nash equilibria tolerant to irrational players. That is, we assume the metric of the given utility functions are accepted by all the parties and then study equilibria robust against (arbitrarily) malicious or envious behaviors of the players. Nevertheless, we believe that the robustness of Nash equilibria studied in this paper can potentially extend to related problems in their setting. Also, while we show general impossibility and computational hardness results of some of our stability notions, it would be interesting to inspect on more specific problems, as those in this line, for the future studies.

Also, the *price of malice*, defined as the ratio between the social welfare when there are malicious players and the social welfare when there is none, was studied by Moscibroda et al. [21] in a virus inoculation game and by Babaioff et al. [4] and Roth [25] in congestion games. However their works are incomparable to ours since we leave out questions about social welfare for now and focus on questions about equilibria themselves.

Rational Cryptography. In contrast to a standard assumption where each player is either honest in follow-

ing the protocol or arbitrarily malicious, rational cryptography concerns scenarios involving both rational and Byzantine players. See [24, 20, 3, 2, 14], to name a few, and surveys by Nielsen [23] and Katz [17]. On one hand, this setting helps the implementation of protocols go beyond the impossibility regarding fairness in the standard setting. On the other hand, it brings new problems in protocol designs that characterize the complex features of player behaviors.

The tolerance threshold of Byzantine players remains an important issue in secure multiparty computation with rational/partial-rational players, and the design of such protocols usually involve the equilibrium properties (instead of social welfare). Therefore the fundamental studies of the existence, complexity, and approximation of the strategy profiles with respect to the stability notions in the paper has a potential for future development in secure protocols under a rich composition of players.

3 Definition

Definition 3.1 (Game). *A normal-form game is described by $G = (\{u_i\}_{i \in [m]}, \{A_i\}_{i \in [m]})$ as follows:*

- m indicates that there are m players $\mathcal{P}_1, \dots, \mathcal{P}_m$.
- u_i is the utility function of \mathcal{P}_i .
- A_i is a finite set of \mathcal{P}_i 's actions.

When A_i 's are all the same, we also denote the game by $G = (\{u_i\}_{i \in [m]}, A)$. We use a bold font \mathbf{x} to denote a random variable and w.l.o.g. a distribution over a given space. We will slightly abuse the notation for the expected value of a utility function u by writing $u(\mathbf{x}) = E_{x \sim \mathbf{x}}[u(x)]$.

Definition 3.2 (Strategy Profile). *For a game $G = (\{u_i\}_{i \in [m]}, \{A_i\}_{i \in [m]})$, a pure strategy profile is a vector $x = (x_1, \dots, x_m)$ in $\bigotimes_{j=1}^m A_j$, and a mixed strategy profile (or profile for short) is a random variable $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_m)$ over $\bigotimes_{j=1}^m A_j$.*

We will also denote a strategy profile by $(\mathbf{x}_{-i} : \mathbf{x}'_i)$, where \mathbf{x}_{-i} is a product distribution over $\bigotimes_{j \neq i}^m A_j$, and $(\mathbf{x}_{-i} : \mathbf{x}'_i) = (\mathbf{x}'_i : \mathbf{x}_{-i}) = (\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}'_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_m)$

Definition 3.3 (Nash Equilibrium). *Let $G = (\{u_i\}_{i \in [m]}, \{A_i\}_{i \in [m]})$ be a game. We say a strategy profile \mathbf{x} is a (mixed) Nash equilibrium if and only if:*

$$\forall i, \forall x'_i \in A_i, u_i(x'_i : \mathbf{x}_{-i}) \leq u_i(\mathbf{x}),$$

For two-payer games, we will denote the players by $\mathcal{P}_0, \mathcal{P}_1$ and the game by $G = (\{u_0, u_1\}, \{A_0, A_1\})$. Similarly, a profile $(\mathbf{x}_0, \mathbf{x}_1)$ indicates that \mathcal{P}_0 plays \mathbf{x}_0 , and \mathcal{P}_1 plays \mathbf{x}_1 , while a profile $(\mathbf{x}_b : \mathbf{x}_{\bar{b}})$, the same as $(\mathbf{x}_{\bar{b}} : \mathbf{x}_b)$, means \mathcal{P}_b plays \mathbf{x}_b and $\mathcal{P}_{\bar{b}}$ plays $\mathbf{x}_{\bar{b}}$ for some $b \in \{0, 1\}$.

4 Warm-up for Immunity and Envy-proofness

We first discuss the connection between the three notions: Nash Equilibrium, immunity, and envy-proofness in two-player games. Interestingly, their connection is very clean. Let's formalize and compare the definitions of these notions in general two-player games:

Definition 4.1 (Immunity for 2-Player Games²). *In a two-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$, we say a strategy profile \mathbf{x} is immune if and only if:*

$$\forall b, \forall x'_b \in A_b, u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) \geq u_{\bar{b}}(\mathbf{x}).$$

Definition 4.2 (Envy-proofness for 2-Player Games). *In a two-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$, we say a strategy profile \mathbf{x} is envy-proof if and only if*

$$\begin{aligned} \forall b, \forall x'_b \in A_b, u_b(\mathbf{x}'_b : \mathbf{x}_{\bar{b}}) - u_b(\mathbf{x}) &\leq u_{\bar{b}}(\mathbf{x}'_b : \mathbf{x}_{\bar{b}}) - u_{\bar{b}}(\mathbf{x}), \\ \text{abbreviated as } \forall b, \Delta_b u_b(\mathbf{x}) &\leq \Delta_b u_{\bar{b}}(\mathbf{x}). \end{aligned} \quad (2)$$

Compare envy-proofness with immunity. Immunity is a stability notion against a complete irrational adversary who can be arbitrarily malicious. She would take any action if it can cause a loss of the other player. On the contrary, envy-proofness is against a partial rational player who is envious about the other player's payoff and compare their payoffs all the time. We strengthen that the goal of envy-proofness is not to divide a cake equally or to make everyone happy (which is a too strong requirement in many real cases). Instead, envy-proofness captures the notion that each player has no incentive to change her mind even if she is an envious player.

Compare envy-proofness with Nash equilibria. While a Nash equilibrium says that a rational player \mathcal{P}_b has no incentive to deviate if $u_b(x'_b : \mathbf{x}_{\bar{b}}) \leq u_b(\mathbf{x})$, an envy-proof profile suggests that an envious player has no incentive to deviate if $u_b(x'_b : \mathbf{x}_{\bar{b}}) - u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) \leq u_b(\mathbf{x}) - u_{\bar{b}}(\mathbf{x})$, which replaces the utility function $u_b(\cdot)$ in the Nash with the difference of the utilities, $(u_b - u'_{\bar{b}})(\cdot)$.

Theorem 1. *For a 2-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$, an immune Nash equilibrium is envy-proof.*

Proof. Let \mathbf{x} be an immune Nash equilibrium. $\forall b, \forall x'_b$, a Nash equilibrium states

$$u_b(x'_b : \mathbf{x}_{\bar{b}}) - u_b(\mathbf{x}) \leq 0,$$

and an immune profile suggests

$$u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) - u_{\bar{b}}(\mathbf{x}) \geq 0,$$

which together imply

$$\Delta_b u_b(\mathbf{x}) \leq 0 \leq \Delta_b u_{\bar{b}}(\mathbf{x}).$$

□

5 On (Non)-Existence

We know that every finite game admits at least one Nash equilibrium [22], but how about the immune profile and the envy-proof profile? Also, does every game admit an immune Nash equilibrium or at least an envy-proof Nash equilibrium? Here we answer these questions for two-player games. Particularly, we show that while every two-player game admits at least one Nash equilibrium, one immune profile, and one envy-proof profile, they are not necessarily to cover each other. That is, some games do not admit any envy-proof Nash equilibrium.

²We note that if we adapt the formula of (ϵ, t) -immunity by Gradwohl and Reingold [13] to two-player games without the relaxation ϵ , it would be $u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) \leq u_{\bar{b}}(\mathbf{x})$ instead. However, we insist on our form that captures the idea that a Byzantine player has no incentive to deviate from an immune profile if he cannot decrease his rival player's payoff by doing so.

	movie	shopping
movie	4,4	1,3
shopping	3,1	0,0

Figure 2: A coordination game. (movie, movie) is the only Nash equilibrium while (shopping, shopping) is the only envy-proof strategy profile.

Lemma 1. *Every finite two-player game admits an immune profile.*

Proof. Let $G = (\{u_0, u_1\}, \{A_0, A_1\})$ be a 2-player game. An immune profile \mathbf{x} requires

$$\forall b, \forall x'_b, u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) \geq u_{\bar{b}}(\mathbf{x}),$$

Define a game $G' = (\{-u_1, -u_0\}, \{A_0, A_1\})$. A Nash equilibrium \mathbf{y} of G' satisfies

$$\forall b, \forall y'_b, -u_{\bar{b}}(y'_b : \mathbf{y}_{\bar{b}}) \leq -u_{\bar{b}}(\mathbf{y}).$$

Hence \mathbf{y} must be an immune profile of G . Since there is always a Nash equilibrium in G' , there is also an immune profile in G . \square

Lemma 2. *Every finite two-player game admits an envy-proof profile.*

Proof. Let $G = (\{u_0, u_1\}, \{A_0, A_1\})$ be a 2-player game. From (2), an envy-proof profile \mathbf{x} says

$$\forall b, \Delta_b u_b(\mathbf{x}) \leq \Delta_b u_{\bar{b}}(\mathbf{x})$$

Define $G' = (\{u_0 - u_1, u_1 - u_0\}, \{A_0, A_1\})$. A Nash equilibrium \mathbf{y} of G' says

$$\forall b, \forall y'_b, u_b(y'_b : \mathbf{y}_{\bar{b}}) - u_{\bar{b}}(y'_b : \mathbf{y}_{\bar{b}}) \leq u_b(\mathbf{y}) - u_{\bar{b}}(\mathbf{y}).$$

Hence \mathbf{y} must be an envy-proof profile of G . Since there is always a Nash equilibrium in G' , there is also an envy-proof profile in G . \square

Lemma 3. *There are two-player games which do not admit any envy-proof Nash equilibrium.*

Proof. Consider a two-player coordination game as described in Figure 2. Note that the game admits only one Nash equilibrium (movie, movie) and only one envy-proof strategy profile (shopping, shopping). Hence there is no envy-proof Nash equilibrium in the game. \square

Corollary 2. *There are two-player games which do not admit any immune Nash equilibrium.*

Proof. This is from Theorem 1 and Lemma 3. \square

6 On Efficient Solvability

After the discussion of (non)-existence, we are curious about whether these notions are efficiently solvable. Particularly for finding a Nash equilibrium, we already know the problem is PPAD-complete, but how about finding an immune profile and find an envy-proof profile?

Lemma 4. *Finding an immune profile in a finite two-player game is PPAD-complete.*

Proof sketch: This is similar to the proof of Lemma 1 but argued in an inverse way. Consider any 2-player game $G = (\{u_0, u_1\}, \{A_0, A_1\})$. Define a game $G' = (\{-u_1, -u_0\}, \{A_0, A_1\})$. Then the problem of finding a Nash equilibrium of G is reduced to finding an immune profile of G' . Since finding Nash is PPAD-complete, it is PPAD-complete, too. \square

Lemma 5. *Finding an envy-proof profile in a finite two-player game is polynomial-time solvable.*

Proof sketch: In light of the proof of Lemma 2, finding an envy-proof profile can reduce to finding a Nash equilibrium of a zero-sum game. Since zero-sum two-player game is polynomial-time solvable, so is finding an envy-proof profile. \square

Recall that from Theorem 1, envy-proofness is a potentially weaker notion than immunity through the reduction sense. Here by Lemma 4 and Lemma 5, it is clear to see this point.

7 Envy-proof Nash and Immune Nash

Now we learned that finding a Nash profile and finding an immune profile are PPAD-hard; meanwhile find an envy-proof profile is efficiently solvable. However, how about immune Nash equilibria and envy-proof Nash equilibria? We have also illustrated that neither an immune Nash equilibrium nor an envy-proof Nash equilibrium does always exist in a game. Hence a follow-up question is whether we can determine their existence in a game efficiently, and whether we can find a solution efficiently if there does exist one.

Gilboa and Zemel [10] and Conitze and Sandholm [7] show that it can be NP-complete to determine the existence of a Nash equilibrium with certain simple constraints. Their results however cannot imply the problem here, nor does their construction in the proof work out. Nevertheless, following the line of reduction from SATISFIABILITY, we show that determining the existence of an envy-proof Nash equilibrium is NP-complete.

Theorem 3. *Even in symmetric 2-player games, it is NP-complete to determine whether there is an envy-proof Nash equilibrium.*

We defer the proof to Section 7.1. The idea is to construct a game such that every legal assignment of a Boolean formula corresponds to an envy-proof profile in the game, and every Nash equilibrium, except a default Nash equilibrium, corresponds to a satisfiable assignment of the formula. Meanwhile, the default, always-existing Nash equilibrium serves as an absorbing state in the game but is not envy-proof.

On the other hand, somehow surprisingly, while finding an immune profile is harder than finding an envy-proof profile (by Lemma 4 and Lemma 5), the following theorem suggests that determining the existence of an immune Nash equilibrium is easier than determining the existence of an envy-proof Nash equilibrium, and immune Nash equilibria can be efficiently located given they exist in a game.

Theorem 4. *Let $G = (\{u_0, u_1\}, \{A_0, A_1\})$ be a normal form 2-player game. It is polynomial-time solvable to determine whether there exists an immune Nash equilibrium in G and to output a solution if there exists.*

Proof. By definition, a strategy profile \mathbf{x} is an immune Nash equilibrium if and only if

$$\forall b, \forall x'_b \in A_b, x'_b \in A_{\bar{b}}, u_b(\mathbf{x}_b : x'_b) \geq u_b(\mathbf{x}) \geq u_b(x'_b, \mathbf{x}_{\bar{b}}). \quad (3)$$

Now define two games G_0, G_1 : For all $b, G_b = (\{u_b, -u_b\}, A_b)$. Note \mathbf{x} is a Nash equilibrium of G_0 and G_1 if and only if \mathbf{x} satisfies (3). Hence \mathbf{x} is an immune Nash equilibrium of G if and only if \mathbf{x} is a Nash equilibrium of G_0 and G_1 at the same time. Because both G_0 and G_1 are zero-sum games, finding their Nash equilibria equals solving linear programming LP_0 and LP_1 : For $b = 0$ and 1 , let M_b be the matrix form of u_b . Then,

$$\begin{aligned} LP_b : \max z_b \\ s.t. \quad & \mathbf{x}^T M_b \geq z_b \mathbf{1}^T \\ & \mathbf{x}^T \mathbf{1} = 1 \\ & \mathbf{x} \geq \mathbf{0} \end{aligned} \quad (4)$$

Clearly G admits an immune Nash equilibrium if and only if LP_0 and LP_1 share a common optimum solution. To check whether they do, we solve LP_0 and LP_1 first. Let $v_b = \max z_b, \forall b$. We claim that the common optimum solution space of LP_0 and LP_1 is nonempty if and only if the following inequations are feasible:

$$\forall b, \mathbf{x}^T M_b \geq v_b \mathbf{1}^T \quad (5)$$

Indeed, if \mathbf{x} is a common optimum solution of LP_0 and LP_1 , then certainly it fulfills (5). On the other hand, if \mathbf{x} fulfills (5), it implies a feasible solution to LP_b with value at least v_b for each b . As v_b is the optimum value of LP_b , \mathbf{x} is then an optimum solution to LP_b for each b . Since linear programming is polynomial-time solvable, so is the problem. □

7.1 Proof of Theorem 3

Let ϕ be a given Boolean formula in conjunctive normal form. Let $V = \{x^1, \dots, x^n\}$ be the set of variables, $L = \{+x^1, -x^1, \dots, +x^n, -x^n\}$ the set of corresponding literals, and C the set of clauses in ϕ . Define a function $\text{LitSet} : V \cup C \rightarrow 2^L$ that outputs the set of literals corresponding to a variable or a clause. That is, $\text{LitSet}(x) = \{+x, -x\}$ if $x \in V$ and $\text{LitSet}(x) = \{\ell \mid x \text{ contains } \ell\}$ if $x \in C$. We construct a symmetric 2-player game $G(\phi)$ in normal form as follows. Let f be a default action, $\Sigma = V \cup L \cup C \cup \{f\}$, $\beta = -n$ (or any sufficiently small value), and the utility function be

- $u_0(\ell_0, \ell_1) = u_1(\ell_0, \ell_1) = \begin{cases} n-1, & \ell_0 \neq -\ell_1 \\ \beta, & \ell_0 = -\ell_1 \end{cases}, \forall \ell_0, \ell_1 \in L$
- $u_0(x_0, x_1) = u_1(x_0, x_1) = \beta, \forall x_0, x_1 \in V \cup C$.
- $u_0(x, \ell) = u_1(\ell, x) = \begin{cases} n, & \ell \notin \text{LitSet}(x) \\ 0, & \ell \in \text{LitSet}(x) \end{cases}, \forall x \in V \cup C$
- $u_0(\ell, x) = u_1(x, \ell) = \begin{cases} n-2, & \ell \notin \text{LitSet}(x) \\ 2(n-1), & \ell \in \text{LitSet}(x) \end{cases}, \forall x \in V \cup C$

Table 1: The table shows the game $G(\phi)$, where $\phi = (x^1 \vee -x^2) \wedge (-x^1 \vee x^2)$. By \times we denote (β, β) since β is just a sufficiently small value for both the players to avoid running into the state.

	x^1	x^2	$+x^1$	$-x^1$	$+x^2$	$-x^2$	$(x^1 \vee -x^2)$	$(-x^1 \vee x^2)$	f
x^1	\times	\times	0,2	0,2	2,0	2,0	\times	\times	3,2
x^2	\times	\times	2,0	2,0	0,2	0,2	\times	\times	3,2
$+x^1$	2,0	0,2	1,1	\times	1,1	1,1	2,0	0,2	1,1
$-x^1$	2,0	0,2	\times	1,1	1,1	1,1	0,2	2,0	1,1
$+x^2$	0,2	2,0	1,1	1,1	1,1	\times	0,2	2,0	1,1
$-x^2$	0,2	2,0	1,1	1,1	\times	1,1	2,0	0,2	1,1
$(x^1 \vee -x^2)$	\times	\times	0,2	2,0	2,0	0,2	\times	\times	3,2
$(-x^1 \vee x^2)$	\times	\times	2,0	0,2	0,2	2,0	\times	\times	3,2
f	2,3	2,3	1,1	1,1	1,1	1,1	2,3	2,3	4,4

$$\begin{aligned} \bullet \quad u_0(f, x) = u_1(x, f) &= \begin{cases} n+2, & x = f \\ n, & x \in V \cup C \\ n-1, & x \in L \end{cases} \\ \bullet \quad u_0(x, f) = u_1(f, x) &= \begin{cases} n+1, & x \in V \cup C \\ n-1, & x \in L \end{cases} \end{aligned}$$

Take $\phi = (x^1 \vee -x^2) \wedge (-x^1 \vee x^2)$ for example. The game $G(\phi)$ is showed in Table 1. Note that ϕ has exactly two solutions, either assigning both variables to true or both to false, and the game has exactly 3 equilibria: 1) both player play $\{+x^1, +x^2\}$ uniformly randomly; 2) both player play $\{-x^1, -x^2\}$ uniformly randomly; 3) both players play f . The first two are envy-proof while the last one is not.

We will proof the following claim, which implies Theorem 3.

Claim 1. *If (ℓ^1, \dots, ℓ^n) satisfies ϕ , then there is an envy-proof Nash equilibrium of $G(\phi)$, where both players play $\{\ell^1, \dots, \ell^n\}$ uniformly randomly. The only other Nash equilibrium is for both players to play f and is not envy-proof.*

Proof. Let μ be the uniform distribution over $\{\ell^1, \dots, \ell^n\}$. First we show that the profile (μ, μ) is indeed an envy-proof Nash equilibrium. Because $G(\phi)$ is a symmetric game, w.l.o.g., we fix player \mathcal{P}_0 's strategy to be μ and consider changes in player \mathcal{P}_1 's strategy. Since $u_1(\ell^i, \ell^j) = n-1$ for all $i, j \in [n]$, playing any distribution ν over $\{\ell^1, \dots, \ell^n\}$ gives the same expected utility $u_1(\mu, \nu) = n-1$ ³. Playing the negation of some ℓ^i gives $u_1(\mu, -\ell^i) = \frac{1}{n} \cdot \beta + \frac{n-1}{n} \cdot n < n-1$, playing a variable or a clause x in $V \cup C$ gives $u_1(\mu, x) = \frac{1}{n} \cdot 0 + \frac{(n-1)}{n} \cdot n = n-1$, and playing f gives $u_1(\mu, f) = n-1$. Hence $u_1(\mu, \mathbf{a}) \leq u_1(\mu, \mu)$ for any \mathbf{a} ; (μ, μ) is a Nash equilibrium. On the other hand, note $u_0(\mu, \nu) = u_1(\mu, \nu)$, $u_0(\mu, -\ell^i) = u_1(\mu, -\ell^i)$, $u_0(\mu, x) = \frac{1}{n} \cdot 2(n-1) + \frac{(n-1)}{n} \cdot (n-2) = n-1 = u_1(\mu, x)$, and $u_0(\mu, f) = n-1 = u_1(\mu, f)$, which together imply $\Delta_1 u_1(\mu, \mu) = \Delta_1 u_0(\mu, \mu)$. Hence (μ, μ) is envy-proof.

As for the other Nash equilibrium (f, f) , since $u_b(f, f)$ is strictly larger than any other utility value, clearly it is a Nash equilibrium. However it is not envy-proof because $-1 = u_1(f, x) - u_1(f, f) > u_0(f, x) - u_0(f, f) = -2$ for all $x \in V \cup C$.

³Recall that when μ, ν is a random variable over the action space, the notation $u_1(\mu, \nu)$ denotes the expected utility $E_{\mu, \nu}[u_1(\mu, \nu)]$

Next we argue that there is no other Nash equilibria. If \mathcal{P}_1 plays f , clearly the unique best response for \mathcal{P}_0 is to also play f since $u_b(f, f)$ gives the unique maximum utility. On the other hand, if \mathcal{P}_1 plays a mixed strategy \mathbf{x}_1 over $V \cup C$, and \mathcal{P}_0 plays a mixed strategy \mathbf{x}_2 over $\Sigma - \{f\}$, the social utility $\bar{u}(\mathbf{x}_1, \mathbf{x}_2)$ is at most $2(n - 1)$. If \mathcal{P}_0 's utility is smaller than n , it is better for \mathcal{P}_0 to play f instead (since $u_0(f, x) = n$ for all $x \in V \cup C$). Otherwise, it implies \mathcal{P}_1 's utility is at most $n - 2$, and it is better for \mathcal{P}_1 to play f instead (since $u_1(x, f) \geq n - 1$ for all x). Hence the profile cannot be a Nash equilibrium. The same analysis holds if \mathcal{P}_1 plays a mixed strategy over Σ with nonzero probability to play actions in $V \cup C$ and \mathcal{P}_0 plays any mixed strategy over Σ for the following two reasons: 1) Because $u_0(f, \ell) = u_0(\ell, \ell)$ and $u_1(\ell, f) = u_1(\ell, \ell)$, playing f instead would not decrease this part's portion of utility. 2) Because (f, f) gives the maximum utility, even the mixed strategy of \mathcal{P}_1 involves nonzero probability of f , playing f instead would not decrease this part's portion of utility either, and vice versa. Hence the support of a Nash equilibrium cannot involve actions in $V \cup C$.

So, we only need to check strategy profiles over $L \cup \{f\}$. However, if \mathcal{P}_1 puts nonzero probability on f , then again, playing f is a strictly better strategy for \mathcal{P}_0 since $u_0(f, f)$ is the unique maximum utility and $u_0(f, \ell) = u_0(\ell, \ell)$. Hence except (f, f) the support of any other Nash equilibrium can only involve actions in L .

The remaining argument to exclude any other Nash equilibrium over L is similar to that of [7]. Note that in this case, the expected utility for each player is at most $n - 1$. If for some $x \in V$, \mathcal{P}_0 puts a probability on playing either $+x$ or $-x$ less than $\frac{1}{n}$, then the expected utility for \mathcal{P}_1 to play x is strictly greater than $\frac{1}{n} \cdot 0 + \frac{n-1}{n} \cdot n = n - 1$. Furthermore, if for some $\ell \in L$, \mathcal{P}_0 puts positive probability on ℓ while \mathcal{P}_1 on $-\ell$, their utility would be less than $n - 1$, and switching to f would be a strictly better strategy. Hence, these cannot be Nash equilibrium, and we can assume that for each $x \in V$, exactly one of $\{+x, -x\}$ is played with probability $\frac{1}{n}$, and \mathcal{P}_0 and \mathcal{P}_1 choose the same literal set. Finally, if they play a set of literals whose corresponding assignment does not satisfy a clause c , then playing c instead would give utility n and would be a better strategy for both players. □

8 On Approximation

Here we study the additive approximation of the three stability notions. For example, while an immune profile says that each player's unilateral deviation cannot decrease the rival player's utility, an ϵ -immune profile says that each player's unilateral deviation cannot decrease the rival player's utility more than ϵ (in absolute value). Intuitively, an approximate immune profile is stable against a lazy Byzantine player who attacks only when he can make sufficiently huge loss of his opponent.

Definition 8.1. Let $G = (\{u_0, u_1\}, \{A_0, A_1\})$ be a 2-player game and \mathbf{x} be a strategy profile.

- x is an ϵ -Nash equilibrium if $\forall b, \forall x'_b, u_b(x'_b : \mathbf{x}_{\bar{b}}) \leq u_b(\mathbf{x}) + \epsilon$.
- x is ϵ -Immune if $\forall b, \forall x'_b, u_{\bar{b}}(x'_b : \mathbf{x}_{\bar{b}}) \geq u_{\bar{b}}(\mathbf{x}) - \epsilon$.
- x is ϵ -Envy-proof if $\forall b, \Delta_b u_b(\mathbf{x}) \leq \Delta_b u_{\bar{b}}(\mathbf{x}) + \epsilon$.

From Lemma 4, we learned that finding an immune profile of a two-player game G is PPAD-complete. In fact, the proof of Lemma 4 indicates that finding an immune profile of G is the same as finding a Nash equilibrium of another game G' which switches the negative utilities of the two players. Thus, applying a sub-exponential algorithm for finding an ϵ -Nash equilibrium by Lipton, Markakis and Mehta [19], we immediately get the following lemma.

Lemma 6. For a two-player game $G = (\{u_0, u_1\}, A)$, where $|A| = n$, there is a $n^{O(\frac{\ln n}{\epsilon})}$ algorithm to find an ϵ -Immune profile.

The algorithm from Lipton et al. [19] is based on exhaustively searching all possible k -uniform ϵ -Nash equilibria, where " k -uniform" means that each strategy is a uniform distribution over a multiset of size k . They show, through the probability method, that there always exists a k -uniform ϵ -Nash equilibria, where the players' utility is close to an exact Nash equilibria so the algorithm works.

Now let's look at a strategy profile with two stability characters again. Note that by Definition 8.1, we say a strategy profile is ϵ -envy-proof Nash equilibrium if it is a Nash equilibrium that is ϵ -envy-proof. Similarly, a strategy profile is an ϵ_1 -envy-proof ϵ_2 -Nash equilibrium if it is an ϵ_2 Nash equilibrium that is ϵ_1 -envy-proof. From Theorem 3, we learned that finding an envy-proof Nash equilibrium of a two-player game is NP-complete. So similarly, we resort for potential approximate solutions. However the theorem from [19] does not apply here. The reason is that even if we can find an ϵ -Nash equilibrium and even if the players' utility can also be close to some exact envy-proof Nash equilibrium, it is unnecessary that the ϵ -Nash equilibrium remains envy-proof. Nevertheless by adapting the probability method to the problem of approximate envy-proofness here, it can be proved that there is a sub-exponential algorithm for either outputting ϵ -envy-proof ϵ -Nash equilibrium or deciding non-existence of any envy-proof Nash equilibrium in a game.

In the following theorem, we show a more generalized result that has a 2ϵ -additive approximation gap between deciding the existence and deciding the non-existence of an approximate envy-proof Nash equilibrium in a game.

Theorem 5. Let $G = (\{u_0, u_1\}, A)$ be a two-player game with $|A| = n$, $\epsilon > 0$, and $k = \frac{3 \ln n}{\epsilon^2}$.

1. There is a $O(n^{3k})$ -time algorithm for computing a ϵ' -envy-proof ϵ -Nash equilibrium, where ϵ' is the minimum approximation factor of envy-proofness over all k -uniform ϵ -Nash equilibria.
2. If $\epsilon' > 2\epsilon$, then there is no $(\epsilon' - 2\epsilon)$ -envy-proof Nash equilibrium in G .

8.1 Proof of Theorem 5

We will first prove that if there exists an ϵ_{ef} -envy-proof Nash equilibrium, then there exists a k -uniform $(2\epsilon + \epsilon_{ef})$ -envy-proof ϵ -Nash equilibrium. Then by an exhaustive search, we can check all $\binom{n+k-1}{k}^2$ possible k -uniform ϵ -Nash equilibrium in time $O(n^{3k})$. Suppose after the exhaustive search, we cannot find any k -uniform ϵ -Nash equilibrium that is $(2\epsilon + \epsilon_{ef})$ -envy-proof, then by contradiction, there does not exist any ϵ_{ef} -envy-proof Nash equilibrium in the game.

For simplicity we assume the utilities have values between 0 and 1, and both players have n pure strategies. We note that the result can be generalized without these restrictions.

Lemma 7. If a game $G = (\{u_0, u_1\}, \{a_0^1, \dots, a_0^n\}, \{a_1^1, \dots, a_1^n\})$ admits an ϵ_{ef} -envy-proof Nash equilibrium $\mathbf{x}^* \equiv (\mathbf{x}_0^*, \mathbf{x}_1^*)$, then for every $\epsilon < 1$, for every $k \geq \frac{3 \ln n}{\epsilon^2}$, there exists a k -uniform $(2\epsilon + \epsilon_{ef})$ -envy-proof ϵ -Nash equilibrium.

Proof. First we construct a profile $\hat{\mathbf{x}} \equiv (\hat{\mathbf{x}}_0, \hat{\mathbf{x}}_1)$ by sampling from $(\mathbf{x}_0^*, \mathbf{x}_1^*)$: For $b \in \{0, 1\}$, form multiset S_b by sampling k pure strategies with replacement according to \mathbf{x}_b^* independently. Let $\hat{\mathbf{x}}_b$ be a mixed strategy that assigns $1/k$ probability to each member in S_b .

Second let $GOAL$ be the event when $\hat{\mathbf{x}}$ is a $(2\epsilon + \epsilon_{ef})$ -envy-proof ϵ -Nash equilibrium, and define events⁴

$$\begin{aligned}\pi_{b,i} &= \{u_b(a_b^i : \mathbf{x}_{\mathbf{b}}^*) \leq u_b(\mathbf{x}^*) + \epsilon\}, \forall b \in \{0, 1\}, i \in [n], \text{ and} \\ \theta_{b,i} &= \{u_b(a_b^i : \mathbf{x}_{\mathbf{b}}^*) - u_b(\mathbf{x}^*) \leq u_{\bar{b}}(a_{\bar{b}}^i : \mathbf{x}_{\bar{\mathbf{b}}}^*) - u_{\bar{b}}(\mathbf{x}^*) + 2\epsilon + \epsilon_{ef}\}, \forall b \in \{0, 1\}, i \in [n] \\ \text{so, } GOAL &= \bigcap_{b,i} (\pi_{b,i} \cap \theta_{b,i}).\end{aligned}$$

Third, in order to apply probability arguments, we need to somehow decompose $\pi_{b,i}$ and $\theta_{b,i}$. Define the following events:

$$\begin{aligned}\phi_b &= \{|u_b(\hat{\mathbf{x}}) - u_b(\mathbf{x}^*)| \leq \epsilon/2\}, \forall b \\ \psi_{b,i} &= \{|u_b(a_b^i : \hat{\mathbf{x}}_{\mathbf{b}}) - u_b(a_b^i : \mathbf{x}_{\mathbf{b}}^*)| \leq \epsilon/2\}, \forall b\end{aligned}$$

It can be verified that $\phi_b \cap \psi_{b,i} \subseteq \pi_{b,i}$, $\forall b, i$, and $\bigcap_{b'} (\phi_{b'} \cap \psi_{b',i}) \subseteq \theta_{b,i}$, $\forall b, i$

To analyze the probability of event ϕ_b , we further define

$$\begin{aligned}\phi_{b,1} &= \{|u_b(\hat{\mathbf{x}}_{\mathbf{b}} : \mathbf{x}_{\mathbf{b}}^*) - u_b(\mathbf{x}^*)| \leq \epsilon/4\}, \forall b \\ \phi_{b,2} &= \{|u_b(\hat{\mathbf{x}}) - u_b(\hat{\mathbf{x}}_{\mathbf{b}} : \mathbf{x}_{\mathbf{b}}^*)| \leq \epsilon/4\}, \forall b\end{aligned}$$

Clearly, $\phi_{b,1} \cap \phi_{b,2} \subseteq \phi_b$. Put them together: $\bigcap_b \phi_{b,1} \bigcap_b \phi_{b,2} \bigcap_{b,i} \psi_{b,i} \subseteq GOAL$

Finally we bound the probability of the events. Due to the construction of $\hat{\mathbf{x}}_{\mathbf{b}}$, $u_b(\hat{\mathbf{x}}_{\mathbf{b}} : \mathbf{x}_{\mathbf{b}}^*)$ (resp. $u_b(\hat{\mathbf{x}})$) is like the average of k independent random variables of the same expected values $u_b(\mathbf{x}^*)$ (resp. $u_b(\hat{\mathbf{x}}_{\mathbf{b}} : \mathbf{x}_{\mathbf{b}}^*)$). Hence by Hoeffding bound [15],

$$\Pr[\phi_{b,j}^c] \leq 2e^{-k\epsilon^2/8}, \forall b \in \{0, 1\}, j \in \{1, 2\}$$

Similarly,

$$\Pr[\psi_{b,i}^c] \leq 2e^{-k\epsilon^2/2}, \forall b \in \{0, 1\}, i \in [n]$$

Then by union bound, we have

$$\Pr[GOAL^c] \leq \sum_{b,j} \Pr[\phi_{b,j}^c] + \sum_{b,i} \Pr[\psi_{b,i}^c] \leq 8e^{-k\epsilon^2/8} + 4ne^{-k\epsilon^2/2} < 1$$

Since $GOAL$ happens with a nonzero probability, there exists a k -uniform $(2\epsilon + \epsilon_{ef})$ -envy-proof ϵ -Nash equilibrium. □

9 Multiplayer Games

We leave several extended questions about multiplayer games for the future, but before the end of this article, we would like to give some more points of view about fault tolerance against envious players in multi-player games. We will define coalitional envy-proofness as a natural extension and relaxation of envy-proofness in multi-player games and define a γ -varied game that captures the smooth change of the derivative utilities between players.

⁴Recall that when $\mathbf{x}_{\mathbf{b}}^*$ is a mixed strategy, our notation $u_b(a_b^i : \mathbf{x}_{\mathbf{b}}^*)$ denotes the the expected utility over the randomness of $\mathbf{x}_{\mathbf{b}}^*$, and $(a_b^i : \mathbf{x}_{\mathbf{b}}^*)$ denotes an unordered strategy profile where \mathcal{P}_b plays a_b^i , and $\mathcal{P}_{\bar{b}}$ plays $\mathbf{x}_{\mathbf{b}}^*$.

9.1 Preliminaries for Multiplayer Games

First let us reminisce two fault tolerance notions in multiplayer games given in [13] that generalize approximate Nash equilibria and approximate immune profiles respectively.

Definition 9.1 ((ϵ, t) -immunity[13]). *In a m -player game $G = (\{u_i\}_{i \in [m]}, A)$, a strategy profile $\mathbf{x} = \mathbf{x}_1 \times \cdots \times \mathbf{x}_m$ is (ϵ, t) -immune if for every set $S \subset N$ of size at most t , every $x'_S \in A^{|S|}$, and every $j \notin S$,*

$$u_j(\mathbf{x}_{-S} : x'_S) \geq u_j(\mathbf{x}) - \epsilon$$

Definition 9.2 ((ϵ, t) -coalitional Nash equilibrium[13]). *In a m -player game $G = (\{u_i\}_{i \in [m]}, A)$, a strategy profile $\mathbf{x} = \mathbf{x}_1 \times \cdots \times \mathbf{x}_m$ is (ϵ, t) -coalitional Nash equilibrium if for every set $S \subset N$ of size at most t , every $x'_S \in A^{|S|}$, and every $i \in S$,*

$$u_i(\mathbf{x}_{-S} : x'_S) \leq u_i(\mathbf{x}) + \epsilon$$

According to [13], a profile is said to be (ϵ, t) -immune "if players' expected utilities do not decrease by more than ϵ when any t other players deviate arbitrarily." Nevertheless the formula in their paper is written as $u_i(\mathbf{x}) \geq u_i(\mathbf{x}_{-S} : x'_S) - \epsilon$. We checked that their main results would hold even if the inequality were revised to $u_i(\mathbf{x}_{-S} : x'_S) \geq u_i(\mathbf{x}) - \epsilon$. Since the later one is more natural and matches the described notion, we adopt it in Definition 9.1.

On the other hand, the notion of coalition has a long history in game theory, and similar notions include strong Nash equilibria [1], coalition-proof Nash equilibria [5] and coalition-proof correlated strategies [27], etc. Compared to these notions, however, as in [13], here we focus on the notion of fault tolerance with a bound on the number of corrupted players.

In [11, 13], one main result is about the robustness of Nash equilibria in a λ -continuous game. We note that the game is equivalent to an alternatively stated γ -sensitive game in [18]) for $\gamma = \lambda/m$, where m is the number of players. Essentially each player's utility is quite insensitive to another player's change in the game, which is a relaxation of the insensitive game.

Definition 9.3 (γ -sensitive game[18, 11, 13]). *A game $G = (\{u_i\}_{i \in [m]}, A)$ is said to be γ -sensitive if*

$$\forall i, \forall x \in A^m, \forall x'_i \in A, \forall j \neq i, |u_j(x_{-i} : x_i) - u_j(x)| \leq \gamma.$$

From [13], it was observed that γ -sensitive games are strongly fault-tolerant with respect to immunity and coalition as the following theorem.

Theorem 6 ([13]). *Let G be an m -player γ -sensitive game. Then every Nash equilibrium of G is $(\gamma t, t)$ -immune and $(3\gamma t, t + 1)$ -coalitional.*

9.2 Coalitional Envy-proofness

Now we extend the definition of envy-proofness for the corresponding fault-tolerant notion in multi-player games, which we call (ϵ, t) -coalitional envy-proofness as follows.

Definition 9.4 ((ϵ, t) -coalitional envy-proofness). *In a m -player game $G = (\{u_i\}_{i \in [m]}, A)$, a strategy profile $\mathbf{x} = \mathbf{x}_1 \times \cdots \times \mathbf{x}_m$ is (ϵ, t) -coalitional envy-proof if for every set $S \subset N$ of size at most t , every $x'_S \in A^{|S|}$, every $i \in S$, and every $j \notin S$,*

$$u_i(\mathbf{x}_{-S} : x'_S) - u_i(\mathbf{x}) \leq u_j(\mathbf{x}_{-S} : x'_S) - u_j(\mathbf{x}) + \epsilon$$

In addition, we define a γ -varied game that is a natural relaxation of a (generalized) team game as follows.

Definition 9.5 (γ -varied game). A game $G = (\{u_i\}_{i \in [m]}, A)$ is said to be γ -varied if

$$\forall i, j, k, x, |\Delta_i u_j(x) - \Delta_i u_k(x)| \leq \gamma.$$

Then we have an analog observation for fault-tolerance with respect to envy-proofness in γ -varied games as follows.

Theorem 7. Let G be an m -player γ -varied game. Then every Nash equilibrium of G is $(\gamma t, t)$ -coalitional envy-proof.

Proof sketch: Let \mathbf{x} be a Nash equilibrium of G . For every $S \subset N, |S| \leq t$, let $S \equiv \{k_1, \dots, k_t\}$ and $S_\ell \equiv \{k_1, \dots, k_\ell\}, \forall \ell \in [t]$. Then for every $x'_S \in A^{|S|}$, and every $i \in S, j \notin S$

$$\begin{aligned} & (u_i(\mathbf{x}_S : x'_S) - u_i(\mathbf{x})) - (u_j(\mathbf{x}_S : x'_S) - u_j(\mathbf{x})) \\ &= \left(u_i(\mathbf{x}_{-k_1} : x'_{k_1}) - u_i(\mathbf{x}) + \sum_{\ell=2}^t u_i(\mathbf{x}_{-S_\ell} : x'_{S_\ell}) - u_i(\mathbf{x}_{-S_{\ell-1}} : x'_{S_{\ell-1}}) \right) \\ & - \left(u_j(\mathbf{x}_{-k_1} : x'_{k_1}) - u_j(\mathbf{x}) + \sum_{\ell=2}^t u_j(\mathbf{x}_{-S_\ell} : x'_{S_\ell}) - u_j(\mathbf{x}_{-S_{\ell-1}} : x'_{S_{\ell-1}}) \right) \\ & \leq \gamma t \end{aligned}$$

□

Next, we generalize [Theorem 1](#) for the corresponding fault-tolerant notions in multi-player games.

Theorem 8. Let \mathbf{x} be a Nash equilibrium in a game G . If \mathbf{x} is (ϵ_1, t) -immune and (ϵ_2, t) -coalitional, then \mathbf{x} is $(\epsilon_1 + \epsilon_2, t)$ -coalitional envy-proof.

Proof sketch: For every $S \subset N$ of size $|S| \leq t$, every $x'_S \in A^{|S|}$, and every $i \in S, j \notin S$, by (ϵ_2, t) -coalition,

$$u_i(\mathbf{x}_{-S} : x'_S) \leq u_i(\mathbf{x}) + \epsilon_2,$$

and by (ϵ_1, t) -immunity,

$$u_j(\mathbf{x}_{-S} : x'_S) \geq u_j(\mathbf{x}) + \epsilon_1.$$

Thus,

$$u_i(\mathbf{x}_{-S} : x'_S) - u_i(\mathbf{x}) \leq u_j(\mathbf{x}_{-S} : x'_S) - u_j(\mathbf{x}) + \epsilon_1 + \epsilon_2.$$

□

Then by [Theorem 8](#) and [Theorem 6](#), we have the following corollary.

Corollary 9. Let G be a γ -sensitive game. Then every Nash equilibrium of G is $(4\gamma t, t)$ -coalitional envy-proof.

Another corollary is about the existence of coalitional envy-proof Nash equilibria in anonymous games, in which the utility of each player is a function of his own action and the empirical distribution of the other players' actions. Formally, a game $G = (\{u_i\}_{i \in [m]}, A)$ is anonymous if for all i , for all permutation $\sigma : [m] \setminus \{i\} \rightarrow [m] \setminus \{i\}$, for all $x \in A^m$, $u_i(x) = u_i(x_{\sigma(1)}, \dots, x_{\sigma(i-1)}, x_i, x_{\sigma(i+1)}, \dots, x_{\sigma(m)})$. In [\[13\]](#), Gradwohl and Reingold showed that if an m -player anonymous game has a constant action space for each player, then the game has an ϵ -Nash equilibrium that is (ϵ, t) -immune and $(4\epsilon, t + 1)$ -coalitional for $t = O(\sqrt{m})$. Together with [Theorem 8](#), we have the following observation.

Corollary 10. For every constant $d \in \mathbb{N}, \epsilon \in (0, 1)$, there exists a constant c such that for every anonymous game $G = (\{u_i\}_{i \in [m]}, A)$ with $|A| = d$, for $t = c\sqrt{m}$, G admits a ϵ -Nash equilibrium which is $(5\epsilon, t)$ -coalitional envy-proof.

Acknowledgement

The author would like to thank Ruta Metha for valuable comments and discussion.

References

- [1] R. J. Aumann. Acceptable points in games of perfect information. *Contributions to the Theory of Games*, IV:287–324, 1959.
- [2] P. D. Azar and S. Micali. Rational proofs. In *STOC’12*, 2012.
- [3] P. D. Azar and S. Micali. Super-efficient rational proofs. In *EC’13*, 2013.
- [4] M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion games with malicious players. In *EC’07*, 2007.
- [5] B. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof nash equilibria. *Pacific J. Math.*, 42:1–12, 1987.
- [6] P.-A. Chen. The effects of altruism and spite on games. PhD thesis, University of Southern California, 2011.
- [7] V. Conitze and T. Sandholm. New complexity results about nash equilibria. *Games and Economic Behavior*, 63:621–641, 2008.
- [8] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, 2007.
- [9] G. Gamow and M. Stern. *Puzzle-math*. Viking Press, 1958.
- [10] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1:80–93, 1989.
- [11] R. Gradwohl and O. Reingold. Fault tolerance in large games. In *PODC’08*, 2008.
- [12] R. Gradwohl and O. Reingold. Partial exposure in large games. *Games and Economic Behavior*, 68(2):602–613, 2010.
- [13] R. Gradwohl and O. Reingold. Fault tolerance in large games. *Games and Economic Behavior*, 86:438–457, 2014.
- [14] A. Groce, J. Katz, A. Thiruvengadam, and V. Zikas. Byzantine agreement with a rational adversary. In *ICALP’12*, 2012.
- [15] W. Hoeffding. Probability inequalities for sums of bounded random variable. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [16] E. Kalai. Partial exposure in large games. *Econometrica*, 72(6):1631–1665, 2004.

- [17] J. Katz. Bridging game theory and cryptography: Recent results and future directions. In *TCC'09*, 2009.
- [18] M. Kearns, M. M. Pai, A. Roth, and J. Ullman. Mechanism design in large games - incentives and privacy. In *ITCS'13*, 2013.
- [19] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *EC'03*, 2003.
- [20] S. Micali. Rational and resilient protocols. In *PODC'14*, 2014.
- [21] T. Moscibroda, S. Schmid, and R. Wattenhofer. The price of malice: A gametheoretic framework for malicious behavior in distributed systems. *Internet Mathematics*, 6(1):27–57, 2010.
- [22] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.
- [23] J. B. Nielsen. Summary report on rational cryptographic protocols. In *ECRYPT Report*, 2007.
- [24] S. J. Ong, D. Parkes, A. Rosen, and S. Vadhan. Fairness with an honest minority and a rational majority. In *TCC'09*, 2009.
- [25] A. Roth. The price of malice in linear congestion games. In *WINE'08*, 2008.
- [26] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.
- [27] D. M. J. Wooders. Coalition-proof equilibria. *Games and Economic Behavior*, 17:82–112, 1996.